PROCEEDINGS OF SCIENCE

Hierarchically deflated conjugate residual

Azusa Yamaguchi[*†]
University of Edinburgh
E-mail: ayamaguc@staffmail.ed.ac.uk

Peter A Boyle
University of Edinburgh
E-mail: paboyle@ph.ed.ac.uk

We present a progress report on a new class of multigrid solver algorithm suitable for the solution of 5d chiral fermions such as Domain Wall fermions and the Continued Fraction overlap. Unlike HDCG [1], the algorithm works directly on a nearest neighbour fine operator. The fine operator used is Hermitian indefinite, for example $\Gamma_5 D_{dwf}$, and convergence is achieved with an indefinite matrix solver such as outer iteration based on conjugate residual. As a result coarse space representations of the operator remain nearest neighbour, giving an 8 point stencil rather than the 81 point stencil used in HDCG. It is hoped this may make it viable to recalculate the matrix elements of the little Dirac operator in an HMC evolution.

34th annual International Symposium on Lattice Field Theory
24-30 July 2016
University of Southampton, UK

[*]Speaker.
[†]Funded by Intel Parallel Computing Centre.

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0).

arXiv:1611.06944v1 [hep-lat] 21 Nov 2016

1. Introduction

Despite the development of revolutionary new multilevel solver algorithms for Wilson Fermions [3, 4, 5, 6, 7] lying nearly ten years in the past, the extension of the approaches to all fermion actions remains somewhat piecemeal. The generalisation to improved Wilson (clover) fermions was made rather rapidly[8], and subsequent variations [9, 10, 11] have included more efficient subspace setup.

The extension domain wall fermions[13, 14] has been studied[12] and an approach made to give a substantial acceleration for valence analysis based on the red-black preconditioned squared operator[1]. The stencil for the squared operator contains all points with taxicab norm less than four, giving 321 points. This has the result that approach is unnattractive for gauge evolution code where, even if the subspace quality can be preserved along an HMC trajectory, the reevaluation of the matrix elements of the little Dirac operator on each timestep in the integrator, for O(50) vectors in the subspace requires naively 15000 matrix multiplies.

Even admitting a constraint, such as a minimum block size of 4^4, the squared operator stencil only reduces to 81 points[1]. It is clear that in order to make a practical algorithm for accelerating HMC evolution with domain wall Fermions we must escape the constraint that the algorithm work on the squared operator, and in order to do this we must first understand why to date only solvers making use of the squared operator have been successful for domain wall Fermions.

2. Spectrum of domain wall fermions

The spectrum of the 5d domain wall fermion operator is illustrated in figure 1. The spectrum for an appropriate negative 5d mass completely encircles and violates the *folklore* present in numerical analysis called the *half-plane condition*[16]. There is a fundamental reason for this folklore: in the infinite volume the spectrum will become dense, and the Krylov solver is then being asked to form an (analytic) polynomial approximation to $\frac{1}{z}$ over an open region encircling the pole. It is impossible to reproduce the phase winding around zero with an analytic function and indeed one can show that minimising the mean square error of a fixed radius circle gives zero for all polynomial coefficients.

In the case of Conjugate Gradient on the Normal Equations (CGNE), which is used to date in RBC-UKQCD domain wall Fermion evolution, the multiplication of each eigenvalue by its conjugate in solving

$$M_{pc}^{\dagger} M_{pc} \psi = \eta$$

places the phase behaviour under control and reduces the problem to a real spectrum, albeit with a squared range of eigenvalue magnitudes.

In the discrete spectrum, finite volume case, we can consider a toy models which also illustrate the problem. If the spectrum consists of N eigenvalues $\lambda_k = e^{i2\pi k/N}$ the conjugate gradient will only converge with an N-term polynomial, which can be analytically arrived at by Gaussian elimination for small N.

In this paper, we propose to solve the phase problem using γ_5 Hermiticity, without squaring the operator, leaving the coarse space representation of the operator still nearest neighbour. Since the

sparsity pattern is preserved this will represent the first true multigrid for five dimensional chiral fermions.

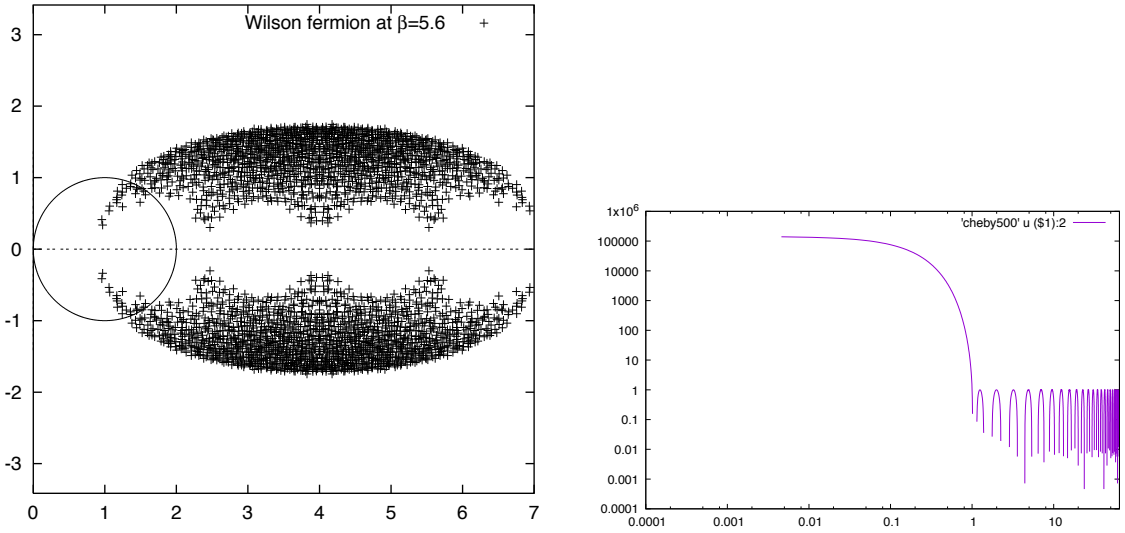

Figure 1: Left: The complex eigenvalue spectrum of the Wilson operator on a 4^4 quenched configuration with $\beta = 5.6$. This figure was produced in reference [15] and we simply reuse the figure here for convenience to illustrate the nature of the spectrum. With a negative Wilson mass of order 1.5, as is introduced in the kernel of chiral fermions, the spectrum is shifted and encircles zero. In five dimensional domain wall and related approaches a five dimensional Wilson term is introduced, without gauge links in the fifth direction, and the "hamburger" picture is repeated with spectrum contained in $0 \leq \Re\lambda \leq 10$ and five "circles". **Right:** an example of an O(500) Chebyshev low-pass filter, shifting the range of the Chebyshev to give exponentially small oscillations around zero between $[\lambda_{\rm cut}, \lambda_{\rm max}]$, with a five order of magnitude enhancement of low modes with $\lambda \leq \lambda_{\rm cut}$.

3. Application to domain wall fermions

We consider two classes of approach for chiral fermions following the nomenclature of ref. [17]. In the Cayley form, the Hermitian indefinite operator for domain wall fermions (and Mobius fermions with $c = 0$, $b \neq 1$) is

$$H_{dwf} = \gamma_5 R_5 D_{dwf} = \Gamma_5 D_{dwf}$$

Meanwhile, the continued fraction form for the standard overlap H_W kernel is already Hermitian indefinite, taking a form that is also appropriate:

$$\begin{bmatrix} H & \frac{1}{\sqrt{\beta_0\beta_1}} & 0 & 0 & 0 \\ \frac{1}{\sqrt{\beta_1\beta_0}} & -H & \frac{1}{\sqrt{\beta_1\beta_2}} & 0 & 0 \\ 0 & \frac{1}{\sqrt{\beta_2\beta_1}} & H & \frac{1}{\sqrt{\beta_2\beta_3}} & 0 \\ 0 & 0 & \frac{1}{\sqrt{\beta_3\beta_2}} & -H & \frac{1}{\sqrt{\beta_3}} \\ 0 & 0 & 0 & \frac{1}{\sqrt{\beta_3}} & R\gamma_5 + \beta_0 H \end{bmatrix}.$$

These operators are *nearest neighbour* and preserve sparsity in a coarse space, but give rise to a Hermitian *indefinite* spectrum. In the infinite volume the spectrum will be dense, real and symmetrical about the origin. From the perspective of a Krylov solver the polynomial approximation $P(\lambda) \sim \frac{1}{\lambda}$ must be made over a the subset real line $\lambda \in [-\lambda_{\max}, -\lambda_{\min}] \cup [\lambda_{\min}, \lambda_{\max}]$

Such a spectrum succumbs easily to the conjugate residual algorithm, which relaxes the Hermitian positive definite constraint of conjugate gradients to only Hermitian indefinite. We will use variants of conjugate residuals as the basis of the outer fine matrix iteration. Regarding the relative efficiency, it is worth to note that we create a Krylov space that strictly contains the CGNE Krylov space (spanned by every second term).

$$P_N(D^\dagger D)D^\dagger = P_N(\Gamma_5 D \Gamma_5 D)\gamma_5 D\gamma_5 \subset P_{2N+1}(H_{dwf})\gamma_5 = P_{2N+1}(\Gamma_5 D)\gamma_5$$

Further, since either on average or in the infinite volume, the spectrum will be symmetrical about zero, the even terms cannot contribute to an approximation of the (odd) function $\frac{1}{x}$ and the in this limit the iteration should converge with an identical number of applications of the nearest neighbour fermion operator as unpreconditioned CGNE. This rule is observed to be almost exactly true even on 16^3 configurations.

4. Two level preconditioner

To introduce a Krylov process as a multigrid preconditioner, we use variable preconditioned GCR as the outer iteration. Since this is a stadard algorithm we do not document it in the interests of brevity. Multigrid may now introduced as the *Preconditioner*. We have tried several approaches to define the low mode vectors used in coarsening. These included i) inverse iteration applied to Gaussian noise, ii) Lanczos eigenvectors, and iii) Chebyshev filters applied to Gaussian noise. An example of our use of high-order Chebyshev filters is given in figure 1. The rapid divergence of a high order Chebyshev outside the default interval $[-1,1]$ is used to *enhance* the modes of interest. We adopt the trick from polynomial preconditioned implicitly restarted Lanczos [18]. Having obtained a basis that captures the near null space of the operator, the vectors are projected into left handed and right handed chiralities. This γ_5 compatible approach was important to eliminate near zero eigenvalues in the coarsened operator[1].

The vectors ϕ_k are then restricted to blocks, of size 2^4 in space time, and the full extent of the fifth dimension, enabling a coarse space representation to be built up as follows.

$$\phi_k^b(x) = \begin{cases} \phi_k(x) & ; x \in b \\ 0 & ; x \notin b \end{cases} \tag{4.1}$$

$$\text{span}\{\phi_k\} \subset \text{span}\{\phi_k^b\}. \tag{4.2}$$

$$P_S = \sum_{k,b} |\phi_k^b\rangle\langle\phi_k^b| \qquad ; \qquad P_{\bar{S}} = 1 - P_S \tag{4.3}$$

$$M = \begin{pmatrix} M_{\bar{S}\bar{S}} & M_{S\bar{S}} \\ M_{\bar{S}S} & M_{SS} \end{pmatrix} = \begin{pmatrix} P_{\bar{S}}MP_{\bar{S}} & P_SMP_{\bar{S}} \\ P_{\bar{S}}MP_S & P_SMP_S \end{pmatrix} \tag{4.4}$$

[1] Suggested to the authors by Kate Clark.

we can represent the matrix M exactly on this subspace by computing its matrix elements, known as the *little Dirac operator* (coarse grid matrix in multi-grid)

$$A^{ab}_{jk} = \langle \phi^a_j | M | \phi^b_k \rangle \qquad ; \qquad (M_{SS}) = A^{ab}_{ij} | \phi^a_i \rangle \langle \phi^b_j |. \tag{4.5}$$

the subspace inverse can be solved by Krylov methods and is:

$$Q = \begin{pmatrix} 0 & 0 \\ 0 & M^{-1}_{SS} \end{pmatrix} \qquad ; \qquad M^{-1}_{SS} = (A^{-1})^{ab}_{ij} | \phi^a_i \rangle \langle \phi^b_j | \tag{4.6}$$

It is important to note that A inherits a sparse structure from M because well separated blocks do *not* connect through M. We can Schur decompose the matrix

$$M = UDL = \begin{bmatrix} M_{\bar{s}\bar{s}} & M_{\bar{s}s} \\ M_{s\bar{s}} & M_{ss} \end{bmatrix} = \begin{bmatrix} 1 & M_{\bar{s}s} M^{-1}_{ss} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} S & 0 \\ 0 & M_{ss} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ M^{-1}_{ss} M_{s\bar{s}} & 1 \end{bmatrix}$$

Note that $P_L M = \begin{bmatrix} S & 0 \\ 0 & 0 \end{bmatrix}$ yields the Schur complement $S = M_{\bar{s}\bar{s}} - M_{\bar{s}s} M^{-1}_{ss} M_{s\bar{s}}$, and that the diagonalisation L and U are projectors P_L and P_R (Galerkin oblique projectors in multi-grid)

$$P_L = P_{\bar{S}} U^{-1} = \begin{pmatrix} 1 & -M_{\bar{S}S} M^{-1}_{SS} \\ 0 & 0 \end{pmatrix} \qquad ; \qquad P_R = L^{-1} P_{\bar{S}} = \begin{pmatrix} 1 & 0 \\ -M^{-1}_{SS} M_{S\bar{S}} & 0 \end{pmatrix} \tag{4.7}$$

We introduce a smoother which is an order 10 Chebyshev polynomial approximation to $1/x$ in the range $[1.0, 64.0]$. To maintain hermiticity in the outer iteration, we presently introduce the smoother and coarse grid preconditioner in a symmetric way, with the composite outer Krylov operating on the matrix as documented in [1]:

$$M_{\text{outer}} = M_{chebyshev} P_L + P_R M_{chebyshev} + Q - M_{chebyshev} P_L \mathscr{H} M_{chebyshev}.$$

5. Initial results

We use a standard RBC-UKQCD $a^{-1} = 1.73$ GeV ensemble with the Iwaski gauge action and DWF 2+1 dynamical flavours with light mass $am_l = 0.01$ and strange mass $am_s = 0.03$ and volume $16^3 \times 32 \times 16$. To make a viable test system, we set the valence mass artificially low to 0.001 to increase the condition number, resulting in thousands of conjugate gradient iterations. We use 16 nodes on Cori phase-1 at NERSC, and take 16 subspace vectors and an order O(900) polynomial.

We display present results from the 16^3 configuration in table 1. A speed up of around a factor of three is obtain in the solution time even form the small volume system. The set up time still presently exceeds the original solve time. The relative speed up is expected to grow as our study progresses to even less well conditioned systems but is the subject of further study.

While it is certainly not yet clear that the final algorithm will be applicable for use in Hybrid Monte Carlo, there are reasons for encouragement. The Lanczos vectors and Chebyshev filtered vectors both demonstrate real speed up over the original red-black conjugate gradient, despite not yet deflating the coarse grid operator. One or two stages of inverse iteration did not yield competitive solution times and appeared less promising as a subspace setup approach. The 40s solve

time was composed of 27s on the fine operator (smoother) and 13s on the coarse space. The coarse space is presently consuming $\frac{1}{3}$ of the time, but has not itself received any further deflation and the algorithm remains strictly two level. Our code implementation in Grid is in principle recursive, and either true recursive multigrid or coarse space eigenvector deflation are open options. Once the coarse space is made cheaper cost can be rebalanced by solving more exactly and using more vectors for the coarse space.

The 300s Chebyshev setup is too long on the test system to be used in HMC on this volume and mass; however since the Chebyshev polynomial is evaluated through a recurrence relation it is also possible to generate Chebyshevs with many different orders for fixed cost. This avenue has not yet been explored. Further, it has become common in multigrid to use polynomial prediction or other schemes to track the subspace across an HMC trajectory since the motion of the gauge configuration field space is limited by the step size, so it is possible this cost could be amortised across a trajectory rather than a single solution.

Algorithm	setup/vecs	Fine Matmuls	Time
CGNE	-	3221	110s
HDCR	Lanczos/16		45s
HDCR	M^{-1}/16		120s
HDCR	Cheby/16		40s
Coarse			13s
Fine		624	27s
Chebyshevs			300s

Table 1: We display the wall clock timing and matrix multiply count for HDCR.

6. Acknowledgements

A.Y. has been supported by an Intel Parallel Computing Centre held at the Higgs Centre for Theoretical Physics. P.B. acknowledges Wolfson Fellowship WM160035, an Alan Turing Fellowship, and STFC grants ST/P000630/1, ST/M006530/1, ST/L000458/1, ST/K005790/1, ST/K005804/1, ST/L000458/1. We would like to thank Martin Luescher, Kate Clark, Evan Weinberg and Richard Brower for useful discussions. All code has been implemented in the Grid library[19], and the Cori phase-1 system at NERSC/LBNL has been used for the code development and testing.

References

[1] P. A. Boyle, "Hierarchically deflated conjugate gradient," arXiv:1402.2585 [hep-lat].

[2] M. Brezina, R. Falgout, S. MacLachlan, T. Manteuffel, S. McCormick, and J. Ruge. Adaptive Smoothed Aggregation (ÎśSA) - SIAM J.Sci.Statist.Comput.,25,1896

[3] M. Luscher, JHEP **0707**, 081 (2007) doi:10.1088/1126-6708/2007/07/081 [arXiv:0706.2298 [hep-lat]].

[4] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn and C. Rebbi, Phys. Rev. Lett. **100**, 041601 (2008) doi:10.1103/PhysRevLett.100.041601 [arXiv:0707.4018 [hep-lat]].

[5] J. Brannick, R. C. Brower, M. A. Clark, J. C. Osborn and C. Rebbi, PoS LAT **2007**, 029 (2007) [arXiv:0710.3612 [hep-lat]].

[6] M. A. Clark, J. Brannick, R. C. Brower, S. F. McCormick, T. A. Manteuffel, J. C. Osborn and C. Rebbi, PoS LATTICE **2008**, 035 (2008) [arXiv:0811.4331 [hep-lat]].

[7] R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen, J. C. Osborn and C. Rebbi, PoS LAT **2009**, 031 (2009) [arXiv:0912.2186 [hep-lat]].

[8] J. C. Osborn, R. Babich, J. Brannick, R. C. Brower, M. A. Clark, S. D. Cohen and C. Rebbi, PoS LATTICE **2010**, 037 (2010) [arXiv:1011.2775 [hep-lat]].

[9] A. Frommer, K. Kahl, S. Krieg, B. Leder and M. Rottmann, PoS LATTICE **2011**, 046 (2011) [arXiv:1202.2462 [hep-lat]].

[10] A. Frommer, K. Kahl, S. Krieg, B. Leder and M. Rottmann, SIAM J. Sci. Comput. **36**, A1581 (2014) doi:10.1137/130919507 [arXiv:1303.1377 [hep-lat]].

[11] A. Frommer, K. Kahl, S. Krieg, B. Leder and M. Rottmann, arXiv:1307.6101 [hep-lat].

[12] S. D. Cohen, R. C. Brower, M. A. Clark and J. C. Osborn, PoS LATTICE **2011**, 030 (2011) [arXiv:1205.2933 [hep-lat]].

[13] D. B. Kaplan, Phys. Lett. B **288**, 342 (1992) doi:10.1016/0370-2693(92)91112-M [hep-lat/9206013].

[14] Y. Shamir, Nucl. Phys. B **406**, 90 (1993) doi:10.1016/0550-3213(93)90162-I [hep-lat/9303005].

[15] W. Bietenholz, hep-lat/0007017.

[16] N. M. Nachtigal, S. C. Reddy, and L. N. Trefethen, "How Fast are Nonsymmetric Matrix Iterations?" SIAM. J. Matrix Anal. Appl., 13(3), 778795.

[17] A. D. Kennedy, hep-lat/0607038.

[18] Rudy Arthur, PhD thesis, University of Edinburgh, 2012.

[19] P. A. Boyle, G. Cossu, A. Yamaguchi and A. Portelli, PoS LATTICE **2015**, 023 (2016).